JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 1968, 11, 651-660 NUMBER 6 (NOVEMBER)

THE REPEATED ACQUISITION OF BEHAVIORAL CHAINS[1]

JOHN J. BOREN AND DENNIS D. DEVINE

WALTER REED ARMY INSTITUTE OF RESEARCH

Monkeys were trained with food reinforcement in a chamber containing four groups of three levers. For each session the monkey's task was to learn a new four-response chain by pressing the correct lever in each group. A stable pattern of learning resulted, and the number of errors reached a steady state from session to session. The technique was then used to determine how various durations of timeouts, following errors, affected the acquisition of new chains. With no timeout, the monkeys made a great many errors, due in large part to superstitious responses within the reinforced chain. Timeout durations ranging from 1 sec to 4 min reduced the number of errors substantially. A second experiment investigated the effects upon acquisition errors of presenting a single light (an "instruction" stimulus) over the correct lever. When this light did not influence the monkeys' responses to the three alternatives, the chains were learned as without it. When the light did control responding, the monkey pressed the appropriate sequence of levers but did not learn the sequence. Thus, when the light was removed, the monkey performed as if learning that sequence for the first time.

CONTENTS

The major goal of the present research was to develop a new technique for the study of a transition phenomenon, specifically the acquisition of chains of behaviors. Behavioral chains were selected for studying acquisition for two reasons. First, many classical studies of learning, such as maze learning or the learning of a sequence of nonsense syllables, have involved chains of behavior. Second, present-day operant research on complex behavior often involves extended behavioral chains. In such studies, the chains are often established by some available method, and then training proceeds until a steady-state baseline is attained. The methods for establishing the baseline behavior are rarely studied in themselves, and the acquisition data often do not appear in the published paper. The present research is a formal elaboration of the beginnings of such steady-state experiments and represents an attempt to explore some of the variables which control the acquisition of behavioral chains.

A second goal of the present research was to develop a technique for studying acquisition with an individual subject. In past studies of behavior acquisition or learning, the experimental design commonly required an independent group of subjects for each value of the independent variable. In contrast, studies of steady-state behavior have often used individual subjects and an experimental design with the following characteristics: (1) each individual subject serves under all experimental conditions; (2) before any variables are manipulated, the subjects are trained until the behavior under study reaches a steady state; and (3) the effect of an independent variable is seen as a change in the steady state. The "individual subject" design has several important advantages over the more conventional "independent groups" design, including elimination of intergroup variability, direct behavioral measures of the individual performance (*versus* statistical derivations), and the direct applicability of the findings to the behavior of the individual (*cf.* Sidman, 1960).

[1]A version of this paper was read at the 1963 meetings of the American Psychological Association. Reprints may be obtained from John J. Boren, Experimental Psychology Dept., Walter Reed Army Institute of Research, Walter Reed Army Medical Center, Washington, D.C. 20012.

In spite of these advantages, acquisition phenomena have not usually been studied with this design. There are several reasons. First, an individual subject's initial acquisition of a performance has often been deemed *the* important learning phenomenon. Since an individual can learn a given problem for the first time only once, each subject can serve under only one experimental condition. Second, as subjects learn such things as new response patterns or new stimulus discriminations, their behavior, far from being at a steady state, is continuously in transition. If a subject learns the same or an equivalent problem a second time, he typically relearns it more quickly and with fewer errors. He is said to exhibit positive transfer. Because of such transfer, a subject's initial learning cannot serve as a control for later learning. Third, without a steady state, there is no stable baseline to be changed by the manipulation of effective variables.

The present techniques were designed to produce a steady state of relearning, thus bypassing the above difficulties and retaining the advantages of the individual subject design. The approach was similar in certain respects to the learning set studies by Harlow (1949), and the behavioral technique was related to the sequential response methods developed by Jenkins (1927) and elaborated more recently by Polidora (1963). A recent paper by Sidman and Rosenberger (1967) is similar both in approach and in technique. In the present study, each subject was trained to acquire different but similar behavioral chains a large number of times so that the pattern of learning and the number of errors reached a steady state from session to session. Then this steady state of repeated acquisition was used as a baseline to study timeouts for errors and a type of stimulus control.

METHOD

Subjects

Three rhesus monkeys served. Their experimental history included more than 100 sessions of training in the chamber described below on response chaining problems. The animals received a maintenance diet of 70 Dietrick and Gambrill lab food pellets (0.7 g) which was earned (except on weekends) during the experimental session. The only additional food was an orange given in the home cage after the session was over. Water was available in the home cage at all times.

Apparatus

The experimental chamber was a metal box measuring 2 ft in each dimension. Twelve levers were mounted in a line on one wall in four groups of three. When a lever was pressed, a relay mounted on the outside wall of the chamber clicked. A pilot light was 2 in. above each lever. Relays, timers, counters, and other electromagnetic equipment controlled and recorded the experiment automatically.

Preliminary Training and General Procedure

Preliminary training involved a number of steps. (1) Each monkey was initially given a food pellet for pressing any of the 12 levers. (2) Next, the reinforcement was contingent upon pressing any one of levers #10, 11, or 12 (the three levers on the monkey's right) when the three pilot lights above these levers were on. (3) Then, the reinforcement was contingent upon a chain of two responses. With the lights on over levers 7, 8, and 9, a press on any of these levers moved the lights on to positions 10, 11, and 12. Then, when the monkey pressed any of levers 10, 11, or 12, it received a pellet. In the same way the chain was gradually extended to include all four groups of levers. (4) The next step was to reinforce the four-member chain on a fixed ratio of two (FR 2); *i.e.*, the pellet was given after the chain was completed twice. This was accomplished by plugging half the holes in a Foringer pellet dispenser. The dispenser sounded after every chain. (5) Each lever in the chain had to be pressed five times (FR 5). (6) Then, when the monkey pressed an inappropriate lever (a lever which was not beneath the pilot lights), the incorrect response was followed by a 15-sec timeout (TO). During TO (or S^{Δ}) both the pilot lights and the chamber illumination were turned off, lever presses were ineffective, and the reinforcement was unavailable. (7) Next, responses on only specified single levers from each set of three were reinforced. For example, levers 2-5-8-11, and only those levers in that sequence, had to be pressed to operate the pellet dispenser. (8) Finally, the correct sequence of lever presses was frequently changed so that the monkey had to learn a new sequence at the beginning of a session.

For example, from session to session the correct sequence was changed at first from 2-5-8-11 to 1-5-8-11 (one lever changed), later from 1-6-8-11 to 2-6-9-11 (two levers changed), and finally from 3-6-9-12 to 2-4-8-10 (all four levers changed).

As a result of the above training, the monkeys eventually acquired the following baseline behavior: at the start of every other session the monkey had to learn which four levers made up the reinforced chain. The first correct lever was always among the three on the monkey's extreme left. The second correct lever was among the next three from the left, and so on as the monkey worked from left to right. When the monkey pressed the correct lever five times, the lights over the next group of three levers came on. In the case of the fourth group of levers, five presses on the correct lever operated the pellet dispenser, and the lights again appeared over the first group of levers. The pellet dispenser operated at the end of every chain, but a pellet was delivered only after every second operation. When the monkey pressed an incorrect lever, the press was immediately followed by a TO (lights out). Any lever press during the TO started the TO interval again. When the TO was over, the pilot lights came back on over the same levers upon which the incorrect response was made, and the monkey could continue to respond for food.

The sequences of lever presses to be reinforced were carefully selected to be equivalent in several ways. First, the sequence was changed from session to session so that a correct lever from the previous session was not repeated. Second, simple orders, such as the first lever in each group of three, were avoided. Third, within a set of six sequences, each lever appeared equally often. An example of a typical set of six sequences is as follows: 3-5-9-11, 2-4-8-10, 1-5-7-11, 2-6-8-12, 3-4-9-10, and 1-6-7-12.

EXPERIMENT I. EFFECTS OF TIMEOUT

This experiment investigated the effect of the TO duration on the frequency of errors. The experiment was carried out after the monkeys had established a stable relearning performance with a TO of 15 sec following each error. Then, TO durations of 1, 15, 60, and 240 sec were compared with no TO. The order of study was as follows: no TO, 1 sec, 15 sec, 240 sec, 60 sec, and no TO. The sequence of required lever presses was changed before every session. A session ended when the monkey had received 70 pellets and thus had performed the chain 140 times. From 8 to 13 sessions were devoted to stabilization on each TO duration.

The results of this experiment were treated quantitatively in terms of the number of times

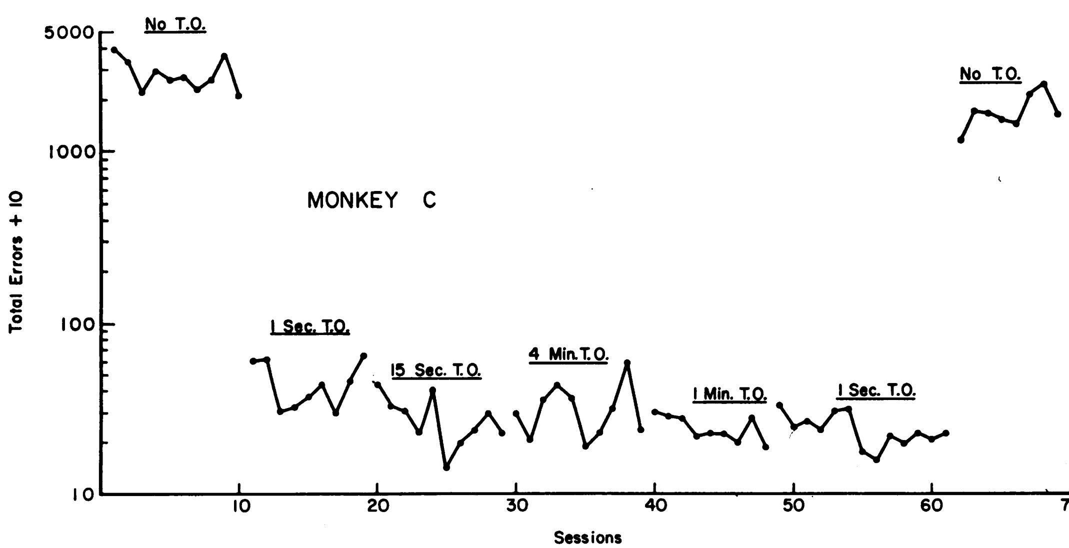

Fig. 1. Total presses on incorrect levers per session under each timeout duration (Monkey C). The ordinate is a logarithmic scale to permit the plotting of a wide range of errors. Ten was added to each error score to avoid the plotting of very small numbers on the log scale. T.O. = timeout.

the monkey pressed a lever which was not part of the reinforced chain. Such lever presses will be referred to as nonreinforced or incorrect responses, or more simply as errors. Figure 1 shows the total number of presses on incorrect levers as Monkey C stabilized on each TO duration. A major point illustrated by this figure is that the acquisition behavior, as measured by the total number of errors under a given TO duration, approached an asymptotic level. Even though the monkey acquired a new behavioral chain each session, the number of errors it made in doing so was reasonably stable. The other monkeys produced data quite similar to those in Fig. 1 and confirmed the major conclusion that the monkeys had established a stable state of relearning. Thus, one objective of this study had been attained. The features of the procedure important in establishing a stable state of relearning seem to be: (1) a set of equivalent problems (lever sequences) to be learned, and (2) a sufficient number of training sessions to produce stability.

The error levels shown in Fig. 1 changed mainly when a TO of any duration was contrasted with no TO. For example, in the first 10 sessions with no TO following incorrect lever presses, the total errors fluctuated around 2000 to 3000 per session. In the next nine ses-

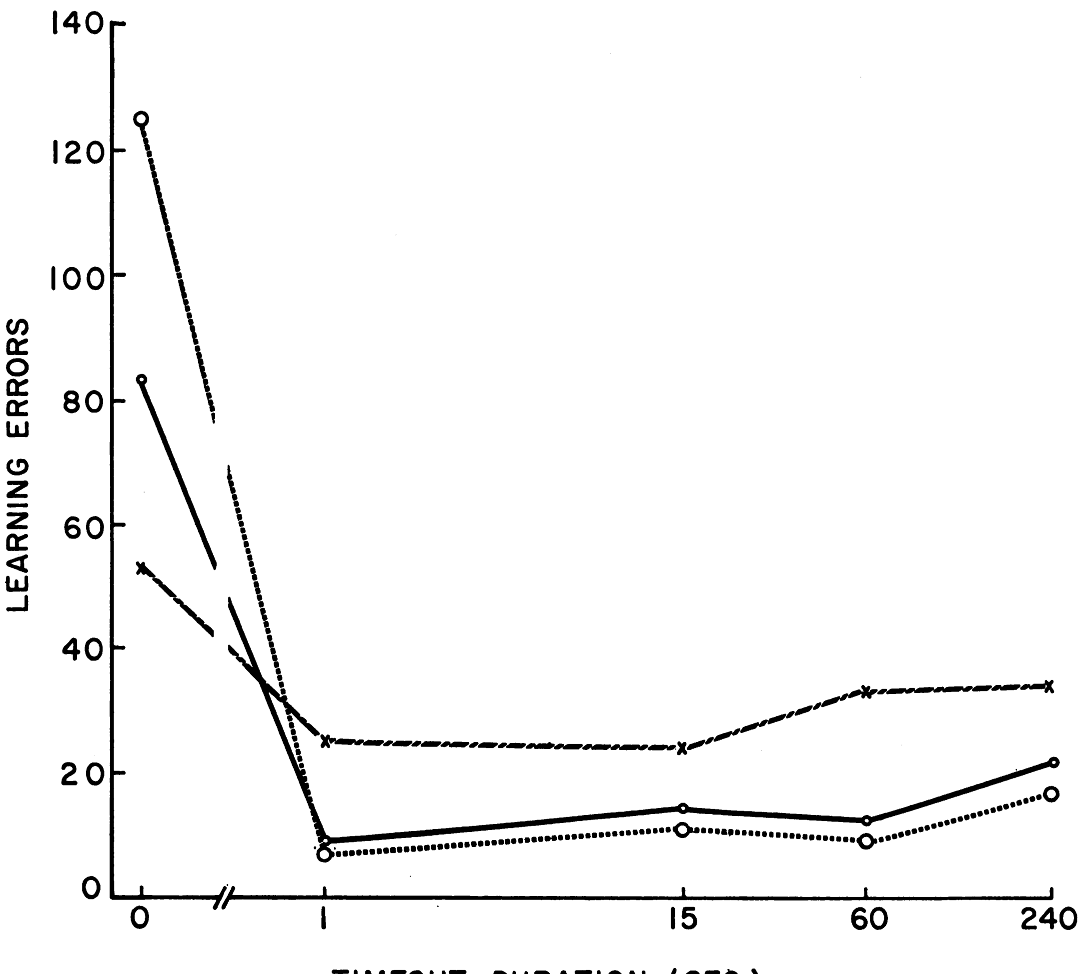

Fig. 2. Mean learning errors as a function of timeout duration. Only the first of a series of presses on an incorrect lever was counted as an error. "Learning" errors were defined as errors occurring during the first five reinforcements of the session. The means were computed over the last six sessions on a given TO duration. The TO durations on the abscissa are on a logarithmic scale. No TO is labeled 0-sec TO. The points with large circles are for Monkey S, with small circles for Monkey C, and X's for Monkey M.

sions, with a TO of only 1 sec following each incorrect lever press, total errors were always less than 60. The surprising finding is that TO durations ranging up to 240 times longer have little, if any, greater effect.

A closer analysis of the errors is provided in Fig. 2 and 3. For these figures, only the first of a series of presses on an incorrect lever was counted. This counting procedure was designed to handle a quantification problem arising from the no-TO procedure. With no TO, a monkey that pressed lever #1 incorrectly might press it 10 or more times before switching to another lever. However, with a TO the monkey's first press on an incorrect lever immediately produced the TO stimulus which terminated any further errors. Thus, the error recording procedure for Fig. 2 and 3 made the errors for the TO and the no-TO conditions more comparable. For Fig. 2 and 3, the errors were also divided into "learning" errors and "performance" errors. "Learning" errors (Fig. 2) refers to the errors during the first five food reinforcements of the session (10 times through a new response chain). This criterion was selected because, when a TO was used, the monkeys had largely acquired the new chain by the tenth repetition. "Performance" errors (Fig. 3) were made during the final 65 reinforcements of the session.

In Fig. 2, the function relating TO duration to learning errors is drawn for each of the

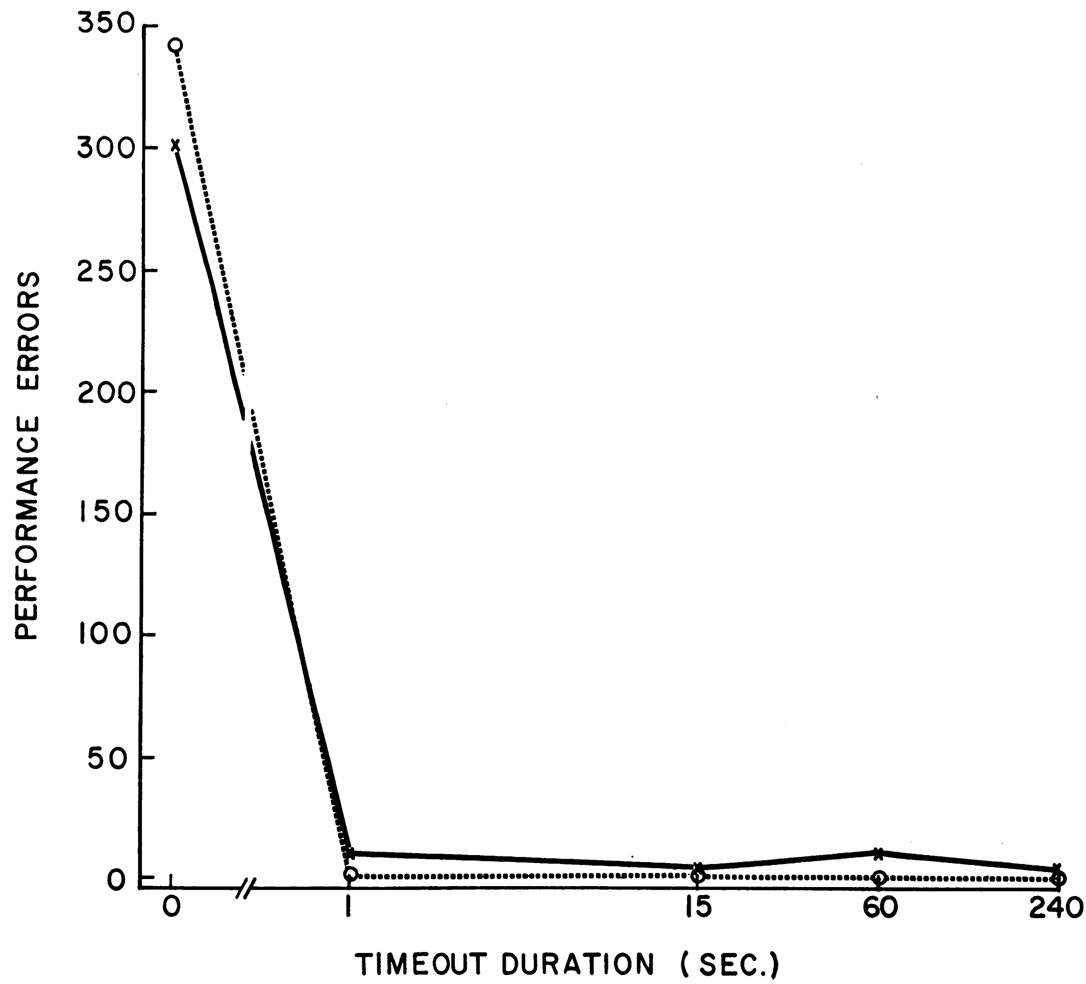

Fig. 3. Mean performance errors as a function of timeout duration. Performance errors were defined as errors occurring during the last 65 reinforcements of the session. As for Fig. 2, the data for Monkey S are plotted with large circles and for Monkey M with X's. The function for Monkey C was almost on top of the functions for Monkeys S and M and therefore is not shown to improve legibility.

three monkeys. The largest number of learning errors occurred when there was no TO (labeled 0-sec TO in the figures). Any duration of TO, from 1 sec to 4 min reduced the number of errors substantially. The function decreased sharply from 0 to 1 sec and then appeared to become asymptotic from 1 to 240 sec.

Experiments with TO duration as the independent variable have been carried out by Ferster and Appel (1961) and by Zimmerman and Ferster (1963). In these studies, pigeons were maintained on a matching-to-sample baseline where incorrect matching responses were followed by a timeout. When matching accuracy was examined as a function of TO duration, it was found that the function passed through a maximum. Intermediate TO's (10 sec to 1 min) produced greater accuracy than either very short TOs (0.5 sec to 1 sec) or very long TOs (2 to 10 min). The function in Fig. 2 for learning errors shows greater effectiveness of a brief TO and less disruption from a long TO. However, a longer TO than was used in the present study might well have produced the disruption observed in the matching-to-sample studies. There was some evidence suggesting this, since one monkey made many errors in a few sessions with the 240-sec TO and failed to complete these sessions.

Figure 3 shows how the number of performance errors during the last 65 reinforcements varied as a function of the TO duration. The function falls sharply between timeout of 0 sec and 1 sec and then levels off from 1 sec through 240 sec. The form of the functional relation in Fig. 3 is much the same as in Fig. 2. However, the absolute number of errors in Fig. 3 differs from Fig. 2. As long as a TO was used, the errors during the last 65 reinforcements were fewer than during the first five reinforcements. However, without a TO, the errors during the last 65 reinforcements were substantially more. Thus, the monkeys seem to have made errors late in the session only when a TO was not used.

Two features of Fig. 2 and 3 require further consideration. (1) Without a TO, the monkeys made a relatively large number of incorrect responses; furthermore, a great many of these incorrect responses were made late in the session. (2) Relatively few incorrect responses occurred with *any* duration of TO from 1 sec to 240 sec. One explanation of these results is that the TO punishes inappropriate behavior. There is good reason to believe that a TO can be aversive and that it can suppress behavior which produced it (Leitenberg, 1965; Azrin and Holz, 1966). Thus, the reduced number of errors found in this study whenever an error produced a TO may be due to punishment by the TOs. However, another mechanism of error reduction may also be involved. The other mechanism was suggested by an analysis of the errors made when no TO was used. The errors were strikingly non-random from chain to chain. Typical examples are illustrated in the event recording on the left of Fig. 4. This record was taken from the middle of a session in which no TO was used and in which the reinforced sequence of levers was 2-4-9-11. This monkey systematically pressed levers 1-2-4-8-9-11 in that order. Responses on levers 1 and 8 were not part of the reinforced chain; yet they intruded repeatedly and consistently. Note, however, that even though the sequence of levers 1-2-4-8-9-11 was needlessly long, it was nevertheless followed by reinforcement. This observation suggests a likely reason for the excess responding: it was probably maintained adventitiously by the normal reinforcements in the situation similar to the superstitious behavior of pigeons observed by Skinner (1948).

The record on the right of Fig. 4 shows the typical behavior when a TO followed any bar-press which was not part of the reinforced chain. The monkey usually pressed only the reinforced levers 2-4-9-11. The superstitious chaining of other lever presses did not occur when a TO followed errors. The TO was related to incorrect responses in the following ways: (1) immediately after the incorrect lever was pressed, the stimulus change marking the TO occurred and further inappropriate behavior was abruptly halted. (2) Then, a delay followed which automatically interposed a temporal space between the incorrect response and the next possible reinforcement. In this way adventitious reinforcement of the inappropriate behavior was averted, and the formation of superstitious chains was made unlikely.

EXPERIMENT II. EFFECTS OF "INSTRUCTIONAL" STIMULI

The second experiment was designed to determine whether a type of stimulus control

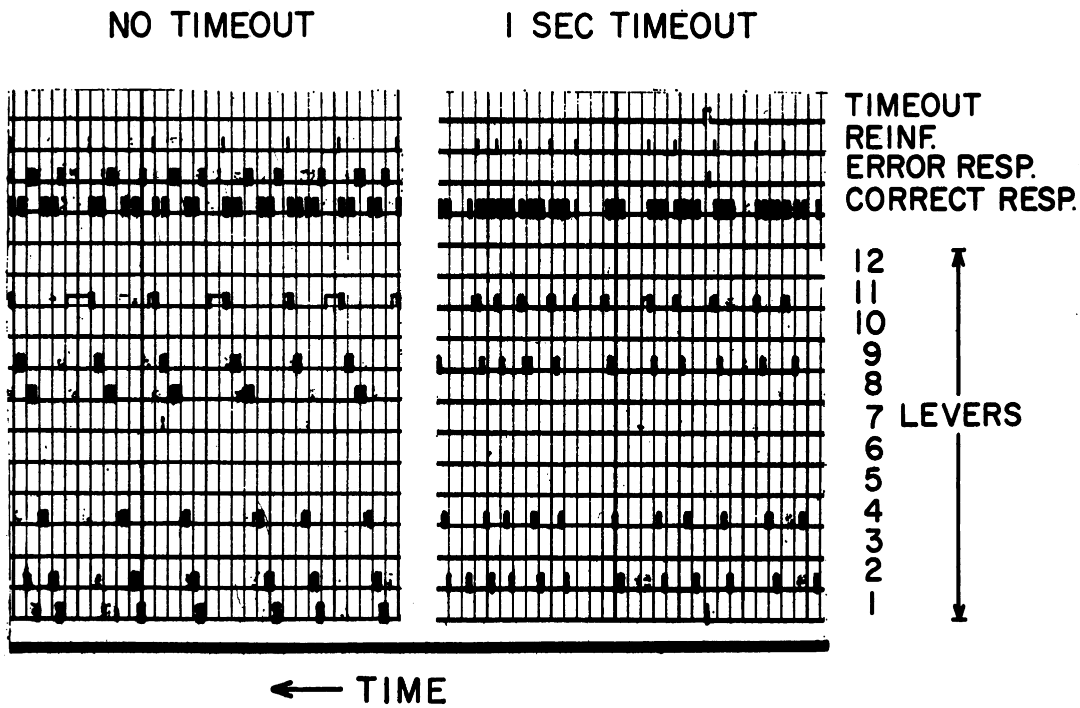

Fig. 4. Event recordings of typical behavior on the chaining procedure. Both records show 1-min segments of Monkey S's performance in the middle of two sessions. The procedure for the panel on the right included a 1-sec timeout after an error. For the panel on the left, a timeout was not used. The reinforced sequence of levers in both records was 2-4-9-11.

would facilitate the acquisition of the behavioral chain. A procedure was arranged which was analogous to instructing a human subject, step by step, exactly what to do. As each new lever sequence was introduced, a single pilot light (an "instruction" stimulus) was turned on over the first correct lever. When the monkey pressed that lever, the first light went out and another light was turned on over the next correct lever, and so on until the entire response sequence had been indicated to the monkey. The experimental question was whether the monkey would learn to make the correct response chain as a result of repeated exposure to this stimulus control procedure, or whether the monkey would merely press the lever under the stimulus without coming under the control of the lever sequence itself.

Several important details of the procedure were changed for Exp. II. Two paired sessions were allotted for the acquisition of each new lever sequence. During the entire first session, a single light was turned on over each correct lever, thus "instructing" the monkey to press the lever below. During the second "non-instructed" session, all three lights were on over the group of three levers, and the experimental issue was whether the monkey would tend to respond correctly as a result of its past history in the previous session.

As a control procedure, another group of paired sessions was run in which the first ("learning") session was a non-instructed, normal session with three lights on over the appropriate three levers. The second ("relearning") session was the same as the first. This control procedure permitted comparison of the facilitative transfer from one normal session to the next *versus* the transfer from an "instruction" session to the next. Blocks of six paired sessions were allocated to the study of both the instruction stimulus procedure and the control procedure. A pair of sessions using instruction stimuli was alternated with a pair using the non-instructed control procedure.

Figure 5 summarizes the complete results for Exp. II. The figure is constructed in the following way. The left half contains the data for Monkey C, the right half for Monkey S. Each bar shows the total errors made in a session while the monkey was learning and performing a certain sequence of lever presses. For example, the first bar in the upper left corner shows that Monkey C made 26 errors learning the first response sequence (levers 3-5-9-11). The next adjacent bar shows that 18 errors were made on the second sequence (2-4-8-10), and so on. Within every grouping of six bars, the required lever sequences were identical.

Consider first the data for Monkey S in the right half of Fig. 5. Compare the two groups of bars in the upper right quadrant showing the errors made without specific instruction stimuli (*i.e.*, when three lights were on over three levers). Monkey S made less than one-quarter as many errors during the second sessions on a sequence (labeled "Relearning") as it made during the initial sessions on that sequence (labeled "Learning"). The average number of errors for the six learning sessions was 191.8 while for the six relearning sessions the average was 45.3. This difference, by the Wilcoxon T test, was statistically significant at the 0.05 level. Thus, without the benefit of specific instruction stimuli, this monkey transferred a substantial part of the performance acquired in the initial learning session over to the second relearning session.

Monkey S's data from the stimulus control procedure are in the quadrant on the lower right in Fig. 5. During the first session on a given sequence, the specific instruction stimuli resulted in a small number of errors. An average of 6.3 errors were made with specific stimuli, compared with 191.8 errors made with non-specific stimuli. It was clear, then, that the single light over the correct lever was an effective stimulus that determined which lever the monkey would press.

The critical information in the experiment emerged from the relearning sessions (with three lights over the group of three levers). A small number of errors in these sessions would indicate that the monkey had acquired the behavioral chain from the instruction stimulus procedure of the prior session. As may be seen in Monkey S's bar graph on the extreme lower right of Fig. 5, this was by no means the outcome. Monkey S made an average of 190.3 errors during these sessions, a value which closely approximates the 191.8 errors made

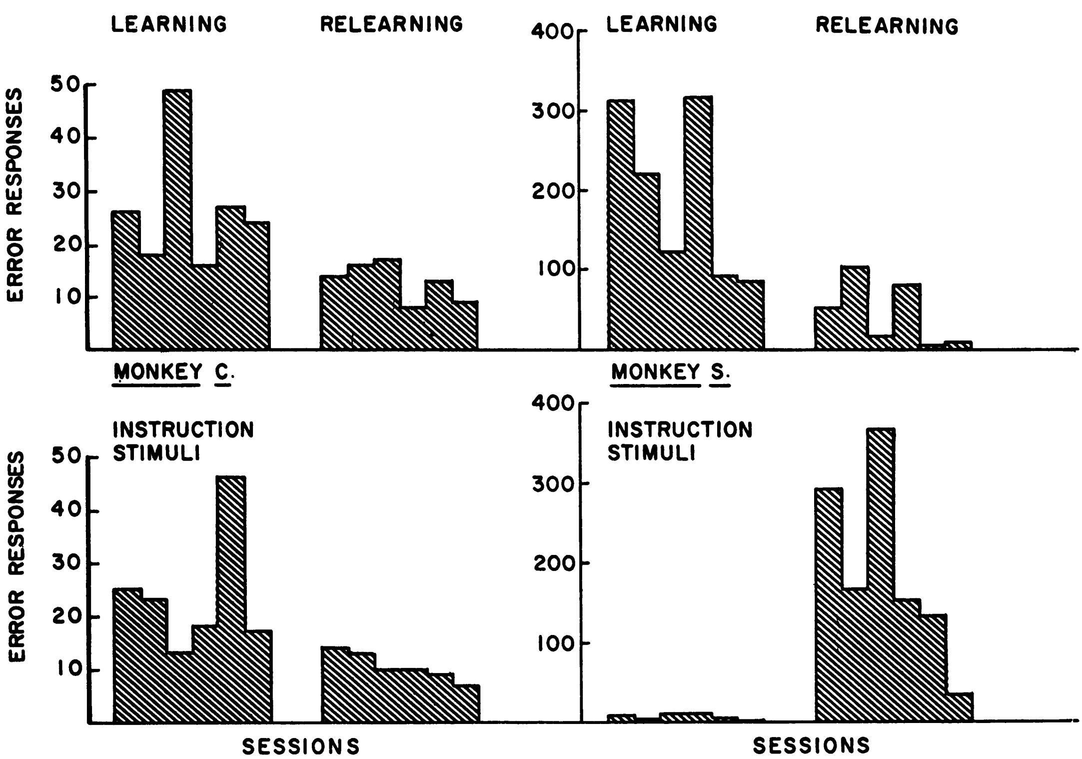

Fig. 5. Number of errors made by Monkeys C and S during the experiment on instruction stimuli. The construction of the figure is explained in the text.

during the initial learning with non-specific stimuli (three lights). The difference was non-significant by the Wilcoxon test. Apparently, then, Monkey S transferred nothing about the response chain from the prior session with specific stimuli, since the error rate was the same as if it were acquiring the chain of lever presses for the first time. The appropriate interpretation seems to be as follows: with instruction stimuli Monkey S merely pressed the lever under the light; thus when this specific stimulus was no longer present, the monkey could not emit the required lever sequence.

Since Monkeys C and M performed quite similarly, the data for only Monkey C are shown on the left half of Fig. 5. The first two groups of bars in the upper left show how many errors Monkey C made without specific stimuli. Less than half as many errors were made during the relearning sessions (12.8 errors) as were made during the learning sessions (26.7 errors). The difference by the Wilcoxon T test was statistically significant at the 0.05 level. Thus, Monkey C (and M) along with Monkey S, show that a substantial amount of the behavior acquired in the initial session transferred over to the second session. Although not presented in Fig. 5, other experiments have shown that the decrease from the second to the third session is relatively small and that the decrease from the third to the fourth session is very small indeed.

The errors made by Monkey C when specific instruction stimuli were used are shown in the quadrant on the lower left of Fig. 5. The results were surprisingly similar to those obtained from this subject with nonspecific stimuli. On the first session of a new lever sequence, Monkey C made an average of 23.7 errors with stimuli and 26.7 errors without stimuli. No statistically significant difference could be shown by the Wilcoxon T test. In keeping with this finding, the number of errors made in the relearning sessions after instructed sessions (mean = 10.5 errors) were essentially the same as after non-instructed sessions (mean the above facts are considered together, one = 12.8 errors). The difference was non-significant by the Wilcoxon T test. When all of the above facts are considered together, one explanation seems to account for the findings: Monkey C was not controlled by the instruction stimuli; or, more technically, the specific stimuli did not determine the specific lever Monkey C pressed. Instead, the specific single light was functionally equivalent to the non-specific three lights. Monkey C was merely pressing one of the three levers in the group beneath the single light in the same way as when three lights were on. Besides the information in Fig. 5, two other sources of information confirm this interpretation. First, direct observation of Monkey C indicated that its eyes were often not oriented toward the instruction light. Second, and more important, misleading instruction stimuli were later arranged in two sessions where the light was *not* over the correct lever. For example, although lever 1 was correct, a light would be on over lever 3; after lever 1 was pressed five times, a light would come on over lever 5, but lever 6 was correct, *etc.* Monkey C's performance was not disrupted by the fraudulent stimuli; instead, it made the usual number of errors. The results for Monkey M were essentially the same. On the other hand, Monkey S made an unusually large number of errors, as would be expected from an animal that had learned to follow instruction stimuli.

This experiment was done twice, with the second version being reported above. Although the first version was marred by minor technical errors, the overall conclusions were identical. Therefore, additional confidence can be placed in the results as described. The reason why Monkeys C and M failed to follow the instruction stimuli is unknown, but it is probably to be found in the 100 or so sessions of preliminary training. In any case, when they entered the experiment, some monkeys were controlled by the instruction stimuli while others were controlled by the lever sequence.

In conclusion, then, the data for Monkeys C and M demonstrate that the instruction stimuli did not change the number of errors made in relearning the four-response chain. The reason in simple terms was that these monkeys failed to follow the instructions. Instead, they learned the sequence of levers to be pressed—which happened to be the exact performance appropriate for the next session. On the other hand, when the instruction stimuli actually controlled the sequence of lever pressing as with Monkey S, the instructions did not improve performance on the second session. Although Monkey S followed the instructions and performed the appropriate response sequence, this monkey did not learn the sequence. In fact, this monkey performed as if it were learning that lever sequence for the first time.

DISCUSSION

The above experiments demonstrate the use of an individual subject, serving as its own control, to study the effects of variables upon a transitional performance. Each individual subject was trained until the transitions had reached a stable level. Then this level could be changed by manipulating relevant variables. The basic technique was to train each subject repeatedly on different but equivalent behavioral chains until the number of acquisition errors decreased to a stable and predictable level. Using this error level as a baseline (analogous to the baseline rate of performance generated by a schedule of reinforcement), the effects of two procedural variables were studied. These variables were the duration of timeouts after errors and the presentation of instruction stimuli. Instruction stimuli were found to have little positive consequence. However, the presentation of a timeout following errors substantially reduced the error level. Errors did not appear to be accidental or uncontrolled. "Correct" behavior in this experiment referred to lever presses upon which the reinforcement was contingent; thus, other lever pressing was defined as "errors". This behavior seemed to be as precisely controlled as the correct behavior. A major cause of errors in this acquisition situation was simply the accidental reinforcement of errors. The errors could be easily chained superstitiously into the correct behavioral sequences. However, the presence of a timeout after errors sharply reduced the superstitious chaining. In addition to these specific results, the study illustrated that the transitional behavior of acquisition could be modified by the same variables in the same way as steady-state operant behavior.

REFERENCES

Azrin, N. H. and Holz, W. C. Punishment. In W. K. Honig (Ed.), *Operant behavior: areas of research and application.* New York: Appleton-Century-Crofts, 1966. Pp. 380-447.

Ferster, C. B. and Appel, J. B. Punishment of S^{Δ} responding in matching to sample by time out from positive reinforcement. *Journal of the Experimental Analysis of Behavior*, 1961, **4**, 45-56.

Harlow, H. F. The formation of learning sets. *Psychological Review*, 1949, **56**, 51-65.

Jenkins, T. N. A standard problem box of multiple complexity for use in comparative studies. *Journal of Comparative Psychology*, 1927, **7**, 129-145.

Leitenberg, H. Is time out from positive reinforcement an aversive event. *Psychological Bulletin*, 1965, **64**, 428-441.

Polidora, V. J. A sequential response method of studying complex behavior in animals and its application to the measurement of drug effects. *Journal of the Experimental Analysis of Behavior*, 1963, **6**, 271-277.

Sidman, M. *Tactics of scientific research.* New York: Basic Books, Inc., 1960.

Sidman, M. and Rosenberger, P. B. Several methods for teaching serial position sequences to monkeys. *Journal of the Experimental Analysis of Behavior*, 1967, **10**, 467-478.

Skinner, B. F. Superstition in the pigeon. *Journal of Experimental Psychology*, 1948, **38**, 168-172.

Zimmerman, J. and Ferster, C. B. Intermittent punishment of S^{Δ} responding in matching to sample. *Journal of the Experimental Analysis of Behavior*, 1963, **6**, 349-356.

Received 3 May 1968.